

Intel in High Performance and Grid Computing

APAC 2005

27 September 2005

Dr. David S. Scott
HPC Technical Director, Asia-Pacific
Intel Technology Asia Pte Ltd
david.s.scott@intel.com



Legal Disclaimer

- Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.
- Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.
- This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.
- Intel® Pentium 4, Prescott, Prestonia, Nocona, Jayhawk, Potomac, and Tulsa processors may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.
- Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at <http://www.intel.com>.
- Copyright © 2004, Intel Corporation. Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.
- *Other brands and names may be claimed as the property of others.



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Agenda

- Intel in HPC
- Intel Software Tools
- Intel in Grid Computing



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Intel's Supercomputer Roots



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.



Process Technology Future

An exponential can't last forever, but you can delay forever for a while

- Gordon Moore



Silicon by the end of the Decade



"30 gigahertz devices, 10 nanometer or less delivering a tera instruction of performance by 2010"⁽¹⁾

1) Pat Gelsinger, Intel CTO, Spring 2002 IDF



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Moore's Law

- Moore's Law is 40 years old!!
- Previous foil is 3 years old
 - Transistor projection still looks good
 - Frequency projection is **NOT**
 - There are several problems, but one of them is that the time to cross the chip doesn't change much.
- How to use the transistors?
- More cache
 - Next Itanium will have 24 MB of cache
- Multicore
 - Dual core everywhere
 - Expect number of cores to follow Moore's Law
- Requires MAJOR shift in software architecture to effectively use many threads



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Technologies

- Multicore (almost) everywhere
- Virtualization Technology
- Power management features
- Active Management Technology
- Reliability features (mapping out bad cache lines)
- Security features (Nx bit and others)
- Fully Buffered DIMMs



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Delivering Greater Performance & Power Efficiency

Example: Server

64-bit Intel® Xeon™ Processor Roadmap

Single Core

1H'06

2H'06

Up to 3.5X
PERFORMANCE/
WATT**

Over 2X
PERFORMANCE**

Intel Next Generation
Micro-architecture
based platform

*Graphics not representative of actual die photo or relative size

With Intel Next Generation Micro-architecture



*Other brands and names are the property of their respective owners

** vs. 64-bit Intel® Xeon™ Processor based platform (as of May '05) *Spec'd rate on DP System

Multi-Core Everywhere

Intel® dual core volume exiting 2006*

Desktop
Performance >70% Dual Core

Server >85% Dual / Multi-Core

Mobile
Performance >70% Dual Core

**17 Multi-Core
CPU projects
(2 in production)**



Desktop Client



Server & Workstation



Mobile client

Intel® Single core to Multi-core shipments



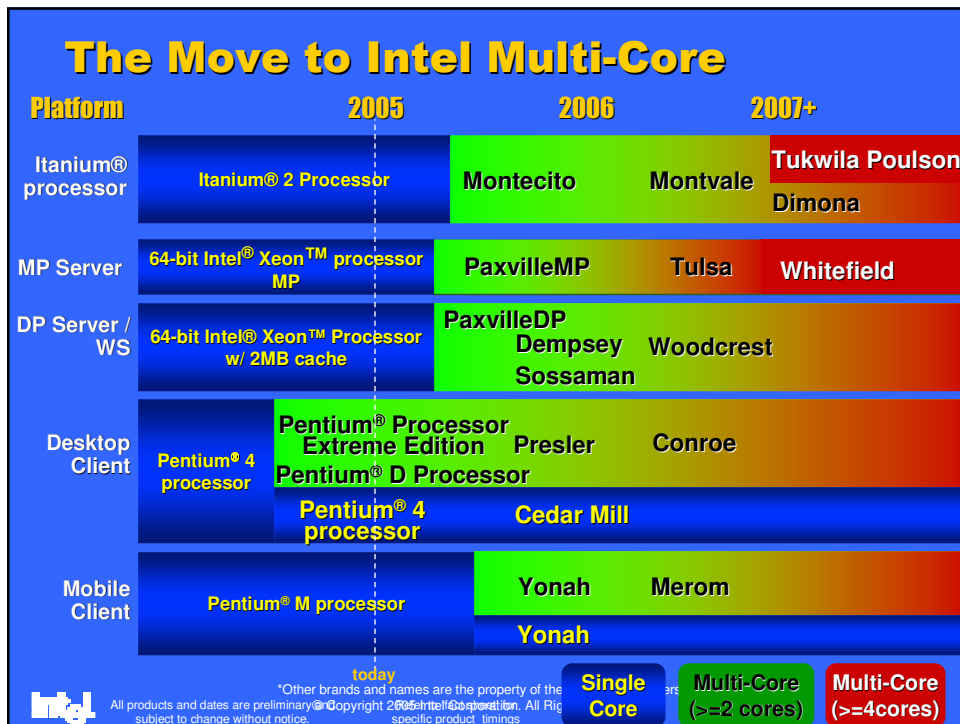
Bringing Parallelism to Everyone



*Other brands and names are the property of their respective owners


© Copyright 2005 Intel Corporation. All Rights Reserved.

*Data is run-rate exiting the year



Intel® Processor based platforms lead the way in the HPC world

- 2/3rd of the top 500 (333) platforms are now based on Intel architecture
 - Twice as prevalent as every other architecture combined.
- Intel Xeon based platforms account for just over half of all of the Top 500 deployments
 - Intel Itanium® based systems are the second most prolific
- >15% year on year growth for Intel architecture
 - 30% year on year growth for Itanium based platforms
 - Vector and RISC based platforms continue to lose ground to COTS based platforms.



List Analysis

- List is very dynamic
- Half life is about 8 months

	Perf	PreviousList
– 500th	1.166TF	299
– 6/05	850Gf	310
– 11/04	624Gf	243
– 6/04	403Gf	290
- 60% of the list are clusters



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

NASA Project Columbia – Out of this world Super Computing

10,240 Itanium processors (20 x 512P)



Climate Predictions Weather Phenomena Re-entry Analysis



Deployed in record-breaking time

Built, installed, and operational on NASA Ames site in < 6 months, ~10X faster over proprietary supercomputers of its size

“The system will be built and integrated over the next 15 weeks (or 3 months). The first two nodes, in fact, were integrated June 28-30 and became operational in early July. This initial build doubled the current capacity at NASA Ames.”

- NASA Press Release July 27, 2004



**Delivering Record breaking computers
In record breaking time**

Intel® Itanium® 2 Processors

Powering World's Fastest Supercomputer Cluster

4096 Itanium processors (4 x 1024)



Record efficiency 86.9%

19.94 teraflops (R_{MAX})

THUNDER

Lawrence Livermore National Laboratory



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Asia-Pacific Systems

- **Generic Supercomputer Centers/National labs**
 - **APAC (1536 Itanium)**
 - Head node of China National grid (1024 Itanium)
 - NCHC in Taiwan
 - KISTI in Korea
- **Mission Labs**
 - BioInformatics Institute in Singapore
 - Central Weather Bureau (Taiwan)
- **Targeted Industrial Applications**
 - Seismic processing (Australia, Malaysia, Indonesia)
 - Digital Content Creation (WETA in NZ, 3256 Xeon)
- **Universities**
 - Hong Kong University
 - Aerospace dept, Seoul National University



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Agenda

- Intel in HPC
- **Intel Software Tools**
- Intel in Grid Computing



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.


Intel® Software Development Products

- Intel® Compilers
 - **Best way to get application performance on Intel processors**
- Intel Performance Libraries
 - **Highly optimized, ready to use building-block functions**
- VTune™ Performance Analyzer
 - **Identify bottlenecks in source code to increase performance or identify problems**
- Intel Threading Tools
 - **Speeds, simplifies development & maintenance of threaded apps**
- Intel Cluster Tools
 - **Create, analyze, optimize and deploy cluster-based applications**
 - **MPI Library “virtualizes” the cluster interconnect**



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

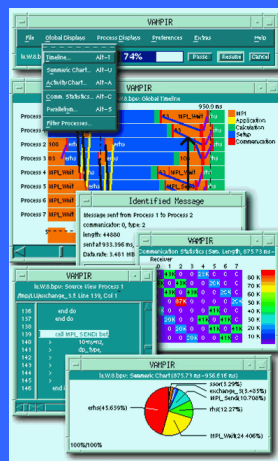
Application Performance



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.


Intel® Cluster Tools

- Intel® Trace Collector
 - Analyzes communication layer and user routines
 - Event Trace based approach
 - Very low impact on application performance
- Intel Trace Analyzer
 - Offline trace analysis for parallel MPI programs
 - High-performance graphics, excellent zooming and filtering
 - Scalability in time and processor-space
- MPI Library
 - Enables application capable of execution across multiple network architectures



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Timelines



*Other brands and names are the property of their respective owners

© Copyright 2005 Intel Corporation. All Rights Reserved. Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States or other countries.

Agenda

- Intel in HPC
- Intel Software Tools
- Intel in Grid Computing



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Value of Grids

- **Ease of use**
 - One login
 - One security check
 - Overview of whole grid
- **Access to**
 - More aggregate compute power available
 - Specialized devices (microscopes, telescopes...)
 - Data anywhere
- **More efficient use of resources**
- **NOT free resources**
- **NOT the world's fastest supercomputer**



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Cluster and Grid Computing

Clusters

- **Homogeneous nodes**
- **Private nodes**
- **Private point-to-point network**
- **Jobs run on collection of nodes**
- **Communication via MPI**

Grids


- **Heterogeneous nodes**
- **Shared nodes**
- **Shared network**
- **Jobs run on one device at a time**
- **Communication by RPC/sockets**



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Designing Chips on an intranet Grid: Netbatch

**To Date,
NetBatch
Has Saved
Intel
Hundreds of
Millions of
Dollars in
Equipment
and Software**



Netbatch has become a fundamental part of our engineering infrastructure with 1.7 to 2 million Netbatch jobs per month

*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.



Successful Intel-based Grid Deployments




*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.



How it all started: UNICORE and UNICORE Plus Projects

- *A production-ready Grid system that connects Supercomputers and Clusters to a Computing Grid*
- Originally developed in German research projects UNICORE (1997-2000) and UNICORE Plus (2000-2003)
- Initiated by leading German Computing Centers
- Evolved to Open Source Grid system that is used world-wide



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

UNICORE goes SourceForge

- UNICORE been released at SourceForge
 - Intel: Client
 - Fujitsu Labs Europe: NJS, Gateway, TSI, UUDB
 - ICM Warsaw: Plug-ins
 - University of Manchester: Resource Broker
 - FZ Jülich: Plug-ins and contributions
- FZ Jülich coordinates activities
 - Intel coordinates Client contributions
- New releases from download pages
- Source Code available from CVS repository
- Support via mailing lists, bug tracker, etc.
- Visit unicore.sourceforge.net



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Why Grid Programming Environment?

Developing Grid Applications too difficult

- Grid Middleware is changing rapidly
- Applications have to be modified for new versions
- Applications for one Grid implementation will not work on others
- Integration in business applications too complex



Need for a high level abstraction
for Grid application programming
based on current standards



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

GPE Roadmap

- **Developer's Release: 1 March 05**
 - Available at UNICORE sourceforge repository
 - Preview
- **Alpha Release: 1 June 05**
 - Packaged version
 - Limited functionality
- **Beta Release: 1 October 05**
 - Full functionality



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

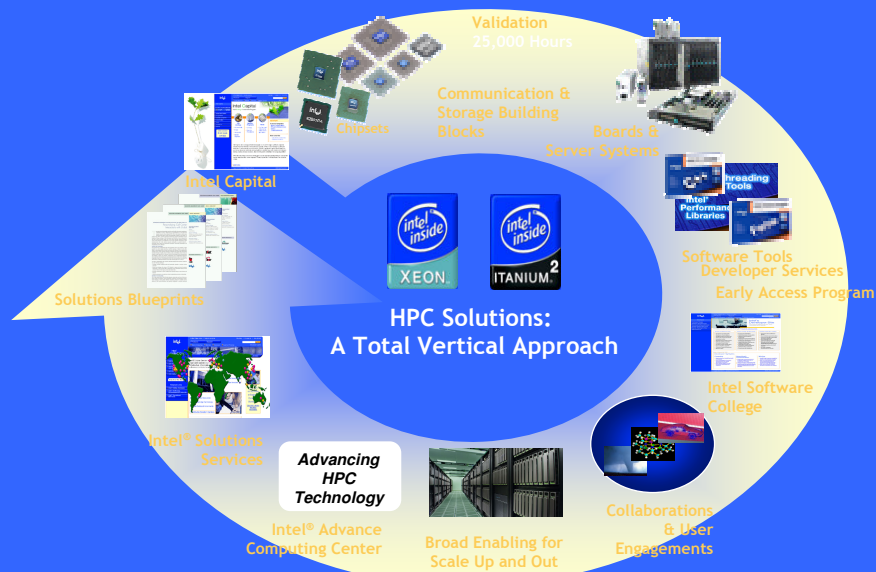
Social & Political Challenges

- Trust establishment
 - by far the single most important challenge
 - Identity or rights
 - Authorized Use
- Education & Support
 - Systems, Network, Security Admins
 - Applications Developers
- Incentives
 - Why would I want to make resources available?
- Organizational Priorities
 - Internal practices may make it hard to use the remote resource
- Managing Expectations
 - Managing the hype.....



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

Not Just Silicon ...Intel has a Deep Commitment to HPC and Grid



*Other brands and names are the property of their respective owners
© Copyright 2005 Intel Corporation. All Rights Reserved.

THANK YOU

